

CLAIMS

What is claimed is:

1. A method for load balancing a plurality of switch modules in a server system, the
5 server system including a plurality of servers, the method comprising the steps of:

a) assigning each of the plurality of servers to a switch module of the plurality
of switch modules, such that a number of servers assigned to each of the plurality of switch
modules is substantially equal.

10 2. The method of claim 1, further comprising the step of:

b) storing for each server a switch assignment in the server's non-volatile storage,
wherein the switch assignment identifies the assigned switch module.

3. The method of claim 2, wherein the assigning step (a) further comprising:

15 (a1) determining whether a server must be assigned to a switch module by
examining the server's non-volatile storage.

4. The method of claim 3, wherein the determining step (a1) further comprising:

(a1i) concluding that the server must be assigned if the non-volatile
20 storage does not include a switch assignment; else

(a1ii) determining whether an existing switch assignment is valid if a
switch assignment is found in the non-volatile storage; and

(a1iii) concluding that the server must be assigned if the existing switch
assignment is invalid.

5. The method of claim 3, wherein the assigning step (a) further comprising:

(a2) accessing a switch assignment table to determine which of the plurality of switch modules to assign to the server if the server must be assigned.

5

6. The method of claim 5, wherein the switch assignment table includes information related to each of the switch modules and the servers assigned thereto.

7. The method of claim 5, wherein the accessing step (a2) further comprising the steps

10 of:

(a2i) utilizing the switch assignment table to evaluate a present load distribution across the plurality of switch modules; and

(a2ii) assigning the server to the next switch module if the number of servers assigned to a current switch module is greater than the number of servers assigned to a next switch module.

15

8. The method of claim 7, wherein the accessing step (a2) further comprising:

(a2iii) assigning the server to the current switch module if the number of servers assigned to the current switch module is not greater than that assigned to the next switch module; else

20

(a2iv) assigning the server to the current switch module if the number of servers assigned to the current switch module is equal to that assigned to the next switch module and the next switch module is a last switch module; else

(a2v) repeating steps (a2ii) through (a2v) if the number of servers assigned to the current switch module is equal to that assigned to the next switch module and the next switch module is not a last switch module, wherein the next switch module becomes the current switch module.

5

9. The method of claim 5, wherein the assigning step (a) further comprising:

(a3) updating the switch assignment table.

10. The method of claim 2 further comprising the step of:

10

(c) utilizing the assigned switch module for access to a network.

11. The method of claim 10, wherein the utilizing step (c) further comprising:

(c1) examining the server's non-volatile storage during a remote boot process to determine whether a switch assignment is present;

15

(c2) requesting, receiving and storing a new switch assignment if the switch assignment is not present; and

(c3) booting via the assigned switch.

12. The method of claim 5, wherein the assigning step (a) is performed during a system

20

power up sequence and when a new server is added to the plurality of servers.

13. A computer readable medium containing program instructions for load balancing a plurality of switch modules in a server system, the server system including a plurality of

servers, the instructions for:

a) assigning each of the plurality of servers to a switch module of the plurality of switch modules, such that a number of servers assigned to each of the plurality of switch modules is substantially equal.

5

14. The computer readable medium of claim 13, further comprising the instruction for:

b) storing for each server a switch assignment in the server's non-volatile storage, wherein the switch assignment identifies the assigned switch module.

10

15. The computer readable medium of claim 14, wherein the assigning instruction (a) further comprising:

(a1) determining whether a server must be assigned to a switch module by examining the server's non-volatile storage.

15

16. The computer readable medium of claim 15, wherein the determining instruction (a1) further comprising:

(a1 i) concluding that the server must be assigned if the non-volatile storage does not include a switch assignment; else

20

(a1 ii) determining whether an existing switch assignment is valid if a switch assignment is found in the non-volatile storage; and

(a1 iii) concluding that the server must be assigned if the existing switch assignment is invalid.

17. The computer readable medium of claim 15, wherein the assigning instruction (a) further comprising:

(a2) accessing a switch assignment table to determine which of the plurality of switch modules to assign to the server if the server must be assigned.

5

18. The method of claim 17, wherein the switch assignment table includes information related to each of the switch modules and the servers assigned thereto.

19. The computer readable medium of claim 17, wherein the accessing instruction (a2) further comprising:

10

(a2i) utilizing the switch assignment table to evaluate a present load distribution across the plurality of switch modules; and

(a2ii) assigning the server to the next switch module if the number of servers assigned to a current switch module is greater than the number of servers assigned to a next switch module.

15

20. The computer readable medium of claim 19, wherein the accessing instruction (a2) further comprising:

(a2iii) assigning the server to the current switch module if the number of servers assigned to the current switch module is not greater than that assigned to the next switch module; else

20

(a2iv) assigning the server to the current switch module if the number of servers assigned to the current switch module is equal to that assigned to the next

switch module and the next switch module is a last switch module; else

(a2v) repeating instructions (a2ii) through (a2v) if the number of servers assigned to the current switch module is equal to that assigned to the next switch module and the next switch module is not a last switch module, wherein the next switch module becomes the current switch module.

21. The computer readable medium of claim 17, wherein the assigning instruction (a) further comprising:

(a3) updating the switch assignment table.

22. The computer readable medium of claim 14 further comprising the instruction for:

(c) utilizing the assigned switch module for access to a network.

23. The computer readable medium of claim 22, wherein the utilizing instruction (c) further comprising:

(c1) examining the server's non-volatile storage during a remote boot process to determine whether a switch assignment is present;

(c2) requesting, receiving and storing a new switch assignment if the switch assignment is not present; and

(c3) booting via the assigned switch.

24. The computer readable medium of claim 17, wherein the assigning instruction (a) is performed during a system power up sequence and when a new server is added to the plurality of

servers.

25. A system for load balancing a plurality of switch modules in a server system, the server system including a plurality of servers, the system comprising:

5 a load balancing mechanism coupled to each of the plurality of servers, wherein the load balancing mechanism assigns each of the plurality of servers to a switch module of the plurality of switch modules, such that a number of servers assigned to each of the plurality of switch modules is substantially equal.

10 26. The system of claim 25, wherein each server comprising:

 means for storing a switch assignment in each server's non-volatile storage, wherein the switch assignment identifies the assigned switch module; and

 means for utilizing the assigned switch module for access to a network.

15 27. The system of claim 26, wherein the load balancing mechanism comprises a switch assignment table, wherein the switch assignment table includes information related to each of the switch modules and the number of servers assigned thereto.

20 28. The system of claim 27, wherein the load balancing mechanism accesses the switch assignment table to determine which of the plurality of switch modules to assign to a server if such an assignment is required.

 29. The system of claim 27, wherein the load balancing mechanism further comprises means for examining a server's non-volatile storage to determine whether the server requires

an assignment to a switch module.

30. The system of claim 27, wherein the load balancing mechanism further comprises means for updating the switch assignment table after a server has been assigned to a switch module.

31. The system of claim 26, wherein the means for utilizing comprising:

means for examining the server's non-volatile storage to determine whether a switch assignment is present,

means for requesting, receiving and storing a new switch assignment if one is not present, and

means for obtaining network access via the assigned switch.

32. A computer system comprising:

a plurality of servers;

a plurality of switch modules coupled to the plurality of servers;

a management module coupled to each of the plurality of servers; and

a load balancing mechanism coupled to the management module, wherein the load balancing mechanism assigns each of the plurality of servers to a switch module of the plurality of switch modules, such that a number of servers assigned to each of the plurality of switch modules is substantially equal.

33. The system of claim 32, wherein each server comprising:

means for storing a switch assignment and a management module identifier in each

server's non-volatile storage, wherein the switch assignment identifies the assigned switch module; and

means for utilizing the assigned switch module for access to a network.

5 34. The system of claim 33, wherein the load balancing mechanism comprises a switch assignment table, wherein the switch assignment table includes information related to each of the switch modules and the servers assigned thereto.

10 35. The system of claim 34, wherein the load balancing mechanism further comprises means for examining a server's non-volatile storage to determine whether the server requires an assignment to a switch module.

15 36. The system of claim 34, wherein the load balancing mechanism further comprises means for updating the switch assignment table after a server has been assigned to a switch module.

37. A method for load balancing a plurality of switch modules in a server system comprising a plurality of servers comprising the steps of:

20 a) assigning each of the plurality of servers to a switch module of the plurality of switch modules, such that a number of servers assigned to each of the plurality of switch modules is substantially equal;

 b) detecting that one of the plurality of servers has been removed; and

 c) if needed, adjusting the load over the plurality of switch modules such that the number of remaining servers assigned to each of the plurality of switch modules is

substantially equal.

38. The method of claim 37, wherein the adjusting step (c) further comprising:

(c1) accessing a switch assignment table, the switch assignment table

including information related to each of the switch modules and the servers assigned thereto;

(c2) comparing the number of remaining servers assigned to each of the plurality of switch modules; and

(c3) based on the comparison, adjusting the load.

39. The method of claim 38, wherein the adjusting step (c3) further comprising:

(c3i) maintaining the load if the number of remaining servers assigned to each of the switch modules is equal, or if the number of remaining servers assigned to each of the switch modules differ by one; else

(c3ii) identifying a switch module having a highest number of servers assigned thereto and choosing a server assigned thereto; and

(c3iii) identifying a switch module having a least number of servers assigned thereto and reassigning the chosen server thereto.

40. A computer readable medium containing program instructions for load balancing a plurality of switch modules in a server system comprising a plurality of servers, instructions for:

a) assigning each of the plurality of servers to a switch module of the plurality

of switch modules, such that a number of servers assigned to each of the plurality of switch modules is substantially equal;

- b) detecting that one of the plurality of servers has been removed; and
- c) if needed, adjusting the load over the plurality of switch modules such that

the number of remaining servers assigned to each of the plurality of switch modules is substantially equal.

41. The computer readable medium of claim 40, wherein the adjusting instruction (c) further comprising:

(c1) accessing a switch assignment table, the switch assignment table including information related to each of the switch modules and the servers assigned thereto;

(c2) comparing the number of remaining servers assigned to each of the plurality of switch modules; and

(c3) based on the comparison, adjusting the load.

42. The computer readable medium of claim 41, wherein the adjusting instruction (c3) further comprising:

(c3i) maintaining the load if the number of remaining servers assigned to each of the switch modules is equal, or if the number of remaining servers assigned to each of the switch modules differ by one; else

(c3ii) identifying a switch module having a highest number of servers assigned thereto and choosing a server assigned thereto; and

(c3iii) identifying a switch module having a least number of servers assigned thereto and reassigning the chosen server thereto.